Patent Application
Docket No. 2003-0225
61922-00010USPT

I hereby certify that this correspondence is being deposited with the U.S. Postal Service as Express Mail, Airbill No. EL258053594US , in an envelope addressed to: MS Patent Application, Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450, on the date shown below.
Dated: 2/5/04 Signature: C Marstaller
(Carol Marstaller)

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICATION FOR U.S. LETTERS PATENT

Title:

**METHOD FOR DETERMINING VOIP GATEWAY PERFORMANCE AND SLAS BASED UPON PATH MEASUREMENTS**

Inventors:

Bob Cole
Howard Lang
Eric Woerner

Steven R. Greenfield
JENKENS & GILCHRIST, A
PROFESSIONAL CORPORATION
1445 Ross Avenue, Suite 3200
Dallas, Texas 75202
214/855-4789

Patent Application
Docket No. 2003-0225
61922-00010USPT

I hereby certify that this correspondence is being deposited with the U.S. Postal Service as Express Mail, Airbill No. EL258053594US , in an envelope addressed to: MS Patent Application, Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450, on the date shown below.

Dated:__________ Signature: __________
(Carol Marstaller)

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICATION FOR U.S. LETTERS PATENT

Title:

# METHOD FOR DETERMINING VOIP GATEWAY PERFORMANCE AND SLAS BASED UPON PATH MEASUREMENTS

Inventors:

Bob Cole
Howard Lang
Eric Woerner

Steven R. Greenfield
JENKENS & GILCHRIST, A
PROFESSIONAL CORPORATION
1445 Ross Avenue, Suite 3200
Dallas, Texas 75202
214/855-4789

TITLE OF INVENTION

# METHOD FOR DETERMINING VOIP GATEWAY PERFORMANCE AND SLAS BASED UPON PATH MEASUREMENTS

## BACKGROUND OF THE INVENTION

Voice-Over-Internet Protocol (VoIP) is attracting a multitude of users because VoIP offers tremendous cost savings relative to a Public Switching Telephone Network (PSTN). For instance, users may bypass long-distance carriers with per minute charges in lieu of transmitting voice calls over the Internet for a flat monthly Internet access fee.

Internet telephony within an intranet enables users to reduce costs by eliminating long-distance charges between sites included in the intranet. An intranet is a local-area network which may or may not be connected to the Internet, but which has some similar functions. The intranet is used for connectivity within, for example, a company. Some companies set up World Wide Web servers on their own internal networks so employees have access to the organization's Web documents. Users may make point-to-point calls via gateway servers attached to a local-area network. For example, a user may want to make a point-to-point call to another user in another office included in the same intranet. The calling party will dial an extension to connect with the gateway server, which is equipped with a telephony board and compression-conversion software; the server configures a private branch exchange (PBX) to digitize the upcoming call. The calling party then dials the number of the called party and the gateway server transmits the call over the IP-based wide-area network to the gateway to the destination office. The destination gateway converts the digital signal back to analog format and delivers the call to the called party.

Although progressing rapidly, VoIP continues to exhibit decreased reliability and sound quality when compared to the PSTN, due primarily to limitations both in Internet bandwidth, current compression technology, delay, jitter, and packet loss. Because the Internet is a packet-switched network, the individual packets of each voice signal may travel over separate network paths for reassembly in the proper sequence at the destination. Although transmitting each packet over a separate path creates a high efficiency for network resources over the PSTN, the chances for packet loss also increase. Packet loss shows up in the form of gaps or periods of silence in a conversation, leading to a clipped-speech effect that is unsatisfactory for most users and unacceptable in business communications. As a result, most corporations looking to reduce communication costs confine their Internet-telephony applications to their intranets. With more predictable bandwidth available than the public Internet, intranets can support full-duplex, real-time voice communications. However, restricting Internet telephony to company intranets does not allow optimum cost saving benefits or flexibility when compared to Internet-telephony over the public Internet.

To date, most developers of Internet-telephony software, as well as vendors of gateway servers, have been using a variety of speech-compression protocols. The use of various speech-coding algorithms, with different bit rates and mechanisms for reconstructing voice packets and handling delays, produces varying levels of intelligibility and fidelity in sound transmitted over the public Internet.

An evolving solution to the varying levels of quality of sound, etc. transmitted over the Internet is to tier the public Internet. Users of the public Internet will then be required to pay for the specific service levels or Quality of Service (QoS) they require. A Service Level Agreement (SLA) is a contract between a carrier and a customer that defines the terms of the carrier's

responsibility to the customer and the type and extent of remuneration if those responsibilities are not met. Reports on the QoS based on either per-call measurements or per-path measurements are a tool for determining if carrier responsibilities are met and if not, any rebate due to the customer. Per-call measurements are capable of illustrating voice quality on a call-by-call basis, which more closely reflects the customer's calling experience. However, in many cases, VoIP gateways or IP-PBXs are managed by the customer and therefore per-call information is not available. For instance, when a customer manages the VoIP gateways, the customer security restrictions or technical constraints may prevent the dissemination of per-call information.

Currently, one solution requires the installation of additional hardware at each customer site to take performance measurements. This solution is not scalable and due to the excessive hardware costs, additional cost of maintaining equipment, and additional network connectivity required to communicate and support the additional hardware, most cost-conscious users would not implement the additional hardware.

Therefore, there is a need for a system and method for making path-based VoIP quality measurements without deploying additional hardware at each customer site.

## SUMMARY OF THE INVENTION

The present invention relates to a method and system for making quality measurements in a VoIP network. More particularly, one aspect of the present invention relates to a system for making quality measurements in a network. The system includes a plurality of routers for routing traffic through the network, means for taking measurements on a path between a first

router and a second router, and means for charging at least one of the plurality of routers when data related to the measurements falls below a target value.

In another aspect, the present invention relates to a method of making quality measurements in a network. The term "R-Factor" as utilized herein refers to an objective measure of voice quality that, for example, accounts for equipment impairments, latency, jitter, and packet loss such as is defined in ITU Standard G.107. The method includes the steps of tracking at least one path that exhibits an R-Factor below a target threshold, tracking a start time indicating when the R-Factor of a particular path falls below the target value, and tracking an end time indicating when the R-Factor of the particular path rises above the target value. The method also includes the steps of determining if an overlap exists between the start time and the end time for multiple paths connecting to a particular router, charging the particular router with one degradation if the overlap exists, and charging the particular router with each degradation if the overlap does not exist.

## BRIEF DESCRIPTION OF THE DRAWINGS

A more complete understanding of the method and apparatus of the present invention may be obtained by reference to the following Detailed Description when taken in conjunction with the accompanying Drawings wherein:

FIGURE 1 is a block diagram of an exemplary VoIP network that may be utilized in accordance with an embodiment of the present invention;

FIGURE 2 is a block diagram of an exemplary VoIP network illustrating multiple routers;

FIGURE 3 is a method for determining the occurrence of degradations in a VoIP network; and

FIGURE 4 is a diagram illustrating a matrix of set and clear events in accordance with an embodiment of the present invention.

## DETAILED DESCRIPTION OF THE INVENTION

Many VoIP customers require SLA support which is achieved by evaluating the quality of calls placed on the network. Most SLAs are related to router performance, not path performance and therefore, embodiments of the present invention translate measurements of the performance of the path between routers into measurements of the performance of the routers. The router performance measurements, in this example, the R-Factor, are used to determine if the QoS guaranteed in the SLA is maintained on a per-site basis. The R-Factor is monitored between designated sites throughout a predetermined amount of time (e.g., a week, a month, etc.) and the QoS guaranteed by the SLA is met if the R-Factor is maintained above a predetermined target value.

Referring now to FIGURE 1, a block diagram of an exemplary VoIP network that may be utilized in accordance with an embodiment of the present invention is illustrated. The VoIP network 100 routes calls from a calling party 102A through a Public Switched Telephone Network (PSTN) 104A to a gateway 106A. As previously mentioned, the gateway 106A may be managed by the customer and therefore, no call detail information is available through the gateway 106A. The call passes through an Ethernet switch 108 to an IP router 110A that may be operated by the supplier of the VoIP service. The call is routed through the Internet 112 to another IP router 110B. The call is then routed through an Ethernet switch 108B, gateway 106B,

and PSTN 104B, to the called party 102B. Embodiments of the present invention are utilized to measure the performance of the path between, for example, IP routers 110A and 110B to determine if the QoS guaranteed is experienced by the calling party 102A and the called party 102B.

Referring now to FIGURE 2, a block diagram of an exemplary VoIP network 100 illustrating a plurality of IP routers 110 is illustrated. Embodiments of the present invention monitor the paths between the customer sites. Therefore, if a problem occurs in a path between routers 110, the particular site that the problem should be attributed to is determined. As evidenced by router A, each router may be a connection point for multiple paths. Therefore, multiple paths may be monitored for a given router 110 such as router A 110A. For router A 110A, the paths between router A 110A and router E 110E, router D 110D, and router C 110C are monitored.

When the R-Factor falls below the target value for a specific path, a degradation is charged to the router 110 at each end of the path. For example, if the path between router A 110A and router C 110C has degraded below the target value, then a degradation is charged to router A 110A and to router C 110C. However, when a router 110 that is a connection point for multiple paths fails, several paths may degrade below the target value. However, in accordance with an aspect of the present invention, if a router 110 that is a connection point for multiple paths fails, the router 110 is charged with one degradation despite the fact that multiple paths may fall below the target value. This aspect prevents double counting of a single failure event. A server 114 responsible for performing path measurements may perform the fore-mentioned calculations, or alternately two separate servers or a separate dedicated device could be utilized.

Referring now to FIGURE 3, a method 300 of implementing an embodiment of the present invention is illustrated. At step 302, at least one path that is below a target value is tracked. At step 304, the start and end times that the path is below the target value is tracked. At step 306 it is determined whether an overlap in the start and end times for a particular router 110 exists. If an overlap exists, at step 308 the particular router 110 is charged for one degradation. If an overlap does not exist, at step 310 the particular router 110 is charged for each degradation occurrence. As evidenced by the exemplary method 300, if a router 110 fails and affects the performance of multiple paths, the start and end times of the failure of each path are tracked. Because the start and end times of each failure will overlap to some extent when the cause is a single failing router 110, the failing router 110 is charged for one degradation despite the fact that multiple degradations are tracked along various paths connected to the failing router. For example, if router 110A is not functioning properly, a below target R-Factor may be seen in paths between routers A 110A and router C 110C, router D 110D, and router E 110E. In this case, router A 110A is charged once, not three times, for falling below the target value.

Referring now to FIGURE 4, an embodiment of the present invention is illustrated as a matrix 400 based on source and destination routers 110. A router pair, such as router A 110A and router B 110B, may experience a problem forcing the R-Factor below the target value. If the R-Factor falls below the target value, then a SET event 402 is written into the matrix 400 at the location of the failure, for example A, B and B, A. When the router pair resumes normal functionality and the R-Factor is restored to a value above the target value, then a CLEAR event 404 is written into the matrix 400 at the same location, in this example, location A, B and B, A.

The matrix 400 may be edited by a manual mechanism (preferably a web interface) that allows the supplier to indicate a site where a problem occurs that results in a breach of the

SLA. In addition, the manual mechanism may also allow the supplier to indicate the nature of the problem (i.e., power failure), a start time indicating when the R-Factor falls below the target value, an end time indicating when the R-Factor rises above the target value, and an identifier (i.e., name, initials, etc.) of the individual that reports the problem. The nature of the problem, as well as any other information may be manually typed in or entered with a drop down menu or other similar data entry means.

When SET events or CLEAR events occur, they may be entered into the matrix 400 in GMT time or the local time where the event occurs. The elapsed time of the SLA breach begins when the first SET event occurs for a particular site, and the elapsed time ends when all the events are CLEARED. It is possible that there are multiple SET events which occur at overlapping times. For the purposes of determining the elapsed time of the SLA breach, the elapsed time begins when the first SET event happens, and ends when the last SET event for that site has CLEARED.

A set of reports may be generated based on the events logged into the matrix 400. For example, an exclusion period report may be generated that lists each site for each customer. The exclusion period report may indicate time periods where the R-Factor is below the target value (if any) along with a reason code that indicates the reason for the R-Factor falling below the target value. The matrix may also be utilized to generate an SLA Report: The SLA report lists each site for each customer and indicates the measured R-Factor for a predetermined time period (e.g., a month), adjusted by any time periods where the R-Factor falls below the target value. The percentage of the month that the measured R-Factor for the site was greater than or equal to the target value may also be shown in the SLA report (the elapsed time that no SET events occur for that site). An R-Factor report may be generated that lists the paths measured for

each customer. The R-Factor report may indicate the percentage of time that the R-Factor is greater than or equal to the target value. The reports may be available for current month-to-date, as well as monthly reports for a prior predetermined amount of time (e.g., 3 months).

Although the above embodiments have been described with reference to intranet/Internet, the present invention may be equally applicable to other environments such as peer-to-peer Internet, solely intranet environments, etc.

The previous description is of various embodiments for implementing the invention, and the scope of the invention should not necessarily be limited solely by these descriptions. The scope of the present invention is instead defined by the following claims.